LA-UR-00-2809

A Geometric Transform for Shape Feature Extraction

Lakshman Prasad, [a] Ramana Rao[b]

[a] Los Alamos National Laboratory, MS E541, Los Alamos, NM, 87545, USA
[b] Integrated Circuit Technology Corp., 2123 Ringwood Ave., San Jose, CA 95131

ABSTRACT

A novel and efficient invertible transform for shape segmentation is defined that serves to localize and extract shape characteristics. This transform—the chordal axis transform (CAT)—remedies the deficiencies of the well-known medial axis transform (MAT). The CAT is applicable to shapes with discretized boundaries without restriction on the sparsity or regularity of the discretization. Using Delaunay triangulations of shape interiors, the CAT induces structural segmentation of shapes into *limb* and *torso* chain complexes of triangles. This enables the localization, extraction, and characterization of the morphological features of shapes. It also yields a pruning scheme for excising morphologically insignificant features and simplifying shape boundaries and descriptions. Furthermore, it enables the explicit characterization and exhaustive enumeration of primary, semantically salient, shape features. Finally, a process to characterize and represent a shape in terms of its morphological features is presented. This results in the migration of a shape from its affine description to an invariant, and semantically salient feature-based representation in the form of attributed planar graphs. The research described here is part of a larger effort aimed at automating image understanding and computer vision tasks.[2, 3, 4, 5, 6]

Keywords: Computer Vision, Delaunay triangulation, feature, morphology, segmentation, shape, skeleton.

1. INTRODUCTION

We present a novel geometric transform, namely the chordal axis transform (CAT), using the constrained Delaunay triangulation (CDT) of polygonized shapes. This creates a paradigm for structural segmentation and analysis of shape data into morphologically meaningful components. The CAT enables the efficient parsing of shapes into semantically significant components that depend only on the shapes' structures and not on their representations per se. In particular, the CAT provides a means of computing with shapes at a high level and understanding them by enabling the transition from a coordinate-based representation of a shape (e.g., by means of pixels or contour points) to a structural representation in terms of its morphological components.

The CAT also provides a means of skeletonization of discretized shapes with the additional facility of prunability to eliminate noisy and morphologically insignificant shape features. This, in particular, has important applications to the automated recognition of printed and handwritten characters and numerals. The skeletal representation of characters of varying thickness and the elimination of insignificant and noisy spurs and branches from the skeleton greatly increases the robustness, reliability, and recognition rates of character-recognition algorithms.

Finally, the CAT paves the way for planar graph-based representation of shapes via their semantically salient shape features. In this form, shapes are highly amenable to comparison and recognition at a semantically higher level of representation.

The method of shape feature extraction described here has direct and immediate applications in several areas such as computer vision, pattern analysis, artificial intelligence, document analysis, optical handwritten character recognition, biometrics, robotic navigation, structural analysis of composite materials,[5] remote surveillance, and content-based image data retrieval.

The notion of "shape" is intimately related to the notion of its *contour* or *boundary*. The boundary of a shape has, however, a continuum of points and as such is not amenable, in general, to finite representation or computation. Thus, in order to characterize a shape computationally, its boundary must be discretized by finitely many points in a morphologically faithful manner, i.e., preserving its structural integrity. Elsewhere,[3] we describe a scale-adaptive multiresolutional scheme using Haar wavelets to discretize shape

contours efficiently. Next, the fundamental morphological attributes of the shape must be extracted from this discrete representation.

The CAT is applicable to shapes with discretized boundaries without restriction on the sparsity of discretization. This is a crucial difference from the MAT, which is not defined for arbitrary discretizations of shapes. It is this property of the CAT that makes it a practically viable and useful construct. The process for obtaining the CAT of a shape with a discretized boundary (hereafter loosely referred to as "discretized shape" or "polygonized shape") involves the construction of the constrained Delaunay triangulation (CDT) of the polygonized shape's interior. The CDT of a polygonal region is not a new construct and has been used elsewhere before for grid generation to numerically solve differential equations with boundary conditions. Here, however, it is used in a novel way to reveal and analyze the semantics and structure of shapes. The CDT serves as a natural morphological grid that localizes the structural properties of the shape, and yields a semantic segmentation of the shape into meaningful components. The CDT effectively localizes morphological events in a shape, such as bifurcations, prolongations, and terminations, and helps identify features, such as protuberances, spans, and handles, that are semantically important hallmarks of visual forms from an anthropocentric and psychophysical perspective.

2. THE CHORDAL AXIS TRANSFORM

When a shape is discretized, its boundary is sampled at discrete points, and neighboring sampled boundary points are joined by straight-line segments. This results in a "polygonal" approximation of the shape (Fig.4). The interior and exterior of the shape are approximated by the interior and exterior of the polygonal approximation.

Before we introduce the Chordal Axis Transform, we will briefly visit the medial axis transform (MAT), and discuss some of its properties and shortcomings.

The *skeleton* (the "frame" over which the "meat" of the shape hangs) is a popular morphological descriptor of a planar shape. For a shape with a continuous boundary, the MAT (Fig.1) yields the generally accepted definition of its skeleton.[1]

Definition 1: The MAT of a planar shape with a continuous boundary is the set of all ordered pairs (p, δ), where p and δ are the center and radius of a *maximal disc* contained in the shape.

A maximal disc contained in the shape is any circle (along with its interior) that is contained completely inside the shape, (i.e., has empty intersection with the exterior of the shape) such that the circle touches the boundary of the shape at two or more points (Fig.1).

The specification of the radii of the maximal discs along with their centers makes the MAT an invertible transform. Indeed, the union of all discs with centers and corresponding radii specified by the MAT of a shape is the shape itself. The medial axis of a shape is just the locus of the centers of all maximal discs contained in the shape. Though the medial axis of a shape has been the generally accepted definition of its skeleton, it is widely acknowledged that the resulting skeleton is not ideal in most cases.

The medial axis of a shape may have several small branches and spurs induced by minor undulations or noise present in the boundary. These features in the medial axis do not contribute significantly to the overall structure of the shape. The MAT of a shape has information only about the distances of the boundary features of the shape and not about their "girths" or "sizes." It is therefore hard to estimate the significance of a point in the MAT to the description of the overall shape. Thus, there is no natural way to "prune" the medial axis or the MAT to obtain a basic skeleton that captures the essence of a shape without regard to the insignificant local boundary features. Although the MAT is defined only for shapes with a continuous boundary, there have been attempts[7] to extend it to shapes with discretized boundaries. This is done by replacing centers and radii of maximal discs by centers and radii of circles that pass through at least three boundary points, and do not contain any boundary points in their interior. However, this extension is unstable under sparse discretizations of the shape boundary, with the medial axis points falling outside the shape (Fig.9). Moreover, the sampling density of the discretization required for the stability of the discrete MAT depends on the local boundary fluctuations and feature sizes, which are usually unknown beforehand.

Further, while it is possible to obtain the MAT of a shape directly from its boundary, it is in general not possible to obtain the boundary of a shape directly by inverting its MAT; the boundary must be recovered by some other means after the shape has been reconstructed. In the case of the discretized MAT, the shape can only be recovered approximately, with the artifacts of reconstruction manifesting as a "bubbly" contour made up of arcs of maximal discs.

The MAT of a small feature on the boundary of a shape can induce a skeletal feature that is spatially far removed from it. In addition, boundary features may be greatly exaggerated or underplayed by their skeletal counterparts. Thus, the MAT of a shape has, in general, not proved to be very useful in providing computationally viable information about the structure of the shape.

A key reason for the practical inadequacy of the MAT is that it constructs the generalized Voronoi Axis of a shape as its descriptor. The Voronoi axis captures the structural ramifications of a shape and provides its medial axis. It is, however, a connected, global object that is sensitive to local perturbations in shape structure. It therefore does not behave gracefully under arbitrary discretizations. It moreover, does not by itself, provide a means of decomposing shape structure into simpler structural components.

For discretized shapes, the CAT overcomes the above problems by considering the geometric dual of the discrete generalized Voronoi axis, namely the constrained Delaunay triangulation (CDT) (Fig.5). The CDT of a polygonized shape decomposes the interior of the shape into triangles, which localize the global structural properties of the shape.

To motivate the choice of the CDT as a tool for shape analysis, we use the idea of constructing maximal discs of shapes as in the case of the MAT. However, in the case of a polygonized shape, circles that pass through at least three vertices of the polygon play the role of maximal discs (Fig.3). These circles have the additional requirement that they do not contain in their interior any vertex of the polygon that is visible to two vertices on the circle. Two vertices u and v of a simple polygon are *visible* to each other, if the line segment joining u and v does not intersect the exterior of the polygon (Fig.2). Such circles are called *empty circles*. For now, we will assume that each empty circle of the polygon has exactly three vertices on it. The case when four or more points are co-circular will be dealt with shortly.

Maximal chords, joining pairs of shape boundary points, are then obtained by joining two nonneighboring vertices of the approximating polygon if and only if an empty circle passes through both these vertices. This results in the CDT of the interior of the polygon (Fig.5). In attempts to discretize the MAT,[7] the centers of empty circles are used to obtain the discrete medial axis, with two centers joined by a line segment if they share a common maximal chord. Indeed, these centers belong to the Voronoi axis of the boundary point set. In general, the center of an empty circle need not lie in the triangle determined by the maximal chords of the circle. This is the source of instability of the discrete MAT under arbitrary discretizations of the shape boundary. Thus, parts of the medial axis skeleton of a polygonized shape may fall outside the shape (Fig.9). To avoid this problem in the case of the skeleton induced by the CAT of a discretized shape, the midpoints of the maximal chords are taken to be the discrete skeletal points (Figs.6, 8, & 12). In an empty circle with two maximal chords, the midpoints of these chords are joined by a line segment. In an empty circle with three maximal chords, the midpoints of the chords are joined to the center of the circle if the center is contained in the triangle, or, are joined to the midpoint of the longest chord otherwise. These constructions are local and internal to the triangle determined by the maximal chords. The result is a connected skeleton that is stable under arbitrary discretizations of shapes, and has the same ramification structure as the medial axis skeleton (Figs.8, 13). The stability of this skeleton is not the only advantage of employing the CDT. As a matter of fact, the skeleton is not the most important descriptor of a shape from the computational point of view. There are many other important advantages to using the CDT as a computational tool for shape analysis. Indeed, in what follows, we will argue that

1) the CDT of a discretized shape contains three kinds of triangles, each characterizing different morphological features,
2) it is possible to identify contiguous chains of triangles in the CDT of a discretized shape that form semantically salient shape features from the point of human visual perception, such as “limbs” and “torsos”,
3) the numbers of each kind of triangle in the CDT of a discretized shape serve to enumerate the salient morphological features of the discretized shape; namely the limb and torso structures,
4) by making appropriate simple, local constructions within each triangle in the CDT of a discretized shape, it is possible to generate a stable skeleton of the overall polygonized shape,
5) the CDT helps simplify and denoise shapes by providing a technique for “pruning” shapes.

We now formally define the CDT of a simple polygon (i.e., a polygon whose boundary does not intersect itself):

Definition 2: The constrained Delaunay triangulation of a simple polygon is a decomposition of a polygon into triangles, such that the circumcircle of each triangle contains no vertex of the polygon inside it that is simultaneously visible to two vertices of the triangle.

In a CDT, the circumcircle of each triangle is an empty circle. However, if an empty circle passes through more than three points, then any triangulation of the interior of the polygon determined by the points on the boundary of the empty circle will be consistent with the above definition. This ambiguity in the CDT for such cases can be resolved meaningfully by additionally constraining the definition of the CDT in the context of shape analysis. To this effect, and in order to define the CAT of polygonized shapes, we must first understand the anatomy of the CDT of shapes.

The triangles of a polygon's CDT can be classified into three types, namely those with two external (i.e., polygonal boundary) edges, those with one external edge, and those with no external edges. Each kind of triangle carries morphological information about the local structure of the polygon. Accordingly, they are given different names. A triangle with two external edges marks the termination of a "limb" or a protrusion of the polygon and is called a *termination triangle* or a ***T****-triangle*. A triangle with one external edge constitutes the "sleeve" of a "limb" or protrusion, signifying the prolongation of the polygon, and is called a *sleeve triangle* or ***S****-triangle*. Finally, a triangle that has no external edges determines a junction or a branching of the polygon, and will accordingly be called a *junction triangle* or a ***J****-triangle*.

In any triangulation of a simple polygon, the number Δ_J of ***J***-triangles is related to the number Δ_T of ***T***-triangles by

$$\Delta_J = \Delta_T + 2g - 2,$$

where g is the number of holes (i.e., the *genus*) of the polygon.

The terms "limb" and "torso" have been loosely and suggestively used above thus far. These terms are formalized below to identify key morphological structures of polygonized shapes. To this effect, two kinds of chain complexes of triangles are identified in any triangulation of a polygon:

Definition 3: A *limb* λ is a chain complex of pairwise adjacent triangles, of the form ***TS . . . SJ*** or ***JS . . . ST*** (Fig.10).

Here, ***J***, ***T***, and ***S*** each symbolize a junction, a terminal, and a sleeve triangle, respectively. The number of sleeve triangles in a limb is allowed to be zero; thus, the duos ***JT*** or ***TJ*** also define limbs.

Definition 4: A *torso* τ is a chain complex of pairwise adjacent triangles, of the form ***JS . . SJ*** (Fig.11).

The ***J***-triangles at the ends of a torso may be the same triangle (as in the case of loops or handles). Again, the number of sleeve triangles in a torso is allowed to be zero; thus, the duo ***JJ*** also defines a torso.

The number of limbs $N(\lambda)$ in any polygon with respect to any triangulation is given by

$$N(\lambda) = \Delta_T / (2 - \min(\Delta_J, 1))$$

and the corresponding number torsos $N(\tau)$ is given by

$$N(\tau) = \tfrac{1}{2}(3\Delta_J - \Delta_T - \min(\Delta_J, 1))$$

We can rewrite the first relation to obtain a formula for the number of holes g in a polygonal shape as well:

$$g = \tfrac{1}{2}(\Delta_J - \Delta_T) + 1$$

Note that these formulae do not involve numbers of ***S***-triangles. This is because ***S***-triangles are not the sites of important morphological events (i.e., bifurcations or terminations); they serve to merely fill up the space between morphological transitions.

For a given polygon, the numbers $N(\lambda)$ and $N(\tau)$ vary, in general, with the triangulation of the polygon, as do the numbers Δ_J and Δ_T. The numbers Δ_J and Δ_T, however, can vary only in accordance with the first formula, involving the genus of a polygon, which is a topological invariant for a given polygon. The formulae for the numbers of limbs and torsos have structural meaning *only* if the triangulation in question is a CDT. In the context of the CDT of a polygon, the above combinatorial formulae have *real* morphological significance. The limb and torso chain complexes of the CDT of a polygon actually *do* correspond to morphological limbs and torsos (i.e., trunks connecting branch points) of the polygon's structure. This is an important and novel property of the CDT of a polygon.

The numbers Δ_J and Δ_T in the CDT of a polygon completely characterize the numbers of all the morphological features of a polygonal shape at the given resolution. These numbers are stable for a shape polygonized at a fixed structural resolution.[3]

We now briefly return to the case of more than three vertices of a polygonized shape lying on an empty circle. As remarked before, any triangulation of the interior of the n-sided polygon $(n>3)$ determined by joining neighboring vertices on the empty circle will be consistent with the definition of the CDT. We can resolve this ambiguity based on the number of internal edges bounding the n-sided polygon to obtain the most desirable triangulation. Each bounding edge of this polygon is either the polygonal boundary edge of the polygonized shape (external edge), or it is an internal edge induced by the CDT (i.e., a maximal chord). Depending on the number of internal bounding edges, we can determine the nature of the triangulation inside the polygon:

Case 1: The polygon has exactly one internal bounding edge.
In this case, the polygon is an n-sided analogue of a terminal triangle, and ought to constitute a limb or the terminal portion of a limb. We therefore choose a triangulation of the polygon that does not contain any junction triangles. Thus, the triangulation will be allowed to have a sequence of ***S***-triangles and one ***T***-triangle. This is done by choosing the vertex furthest from the internal edge and joining its neighboring vertices to form the only ***T***-triangle. The rest of the triangles are constructed to be ***S***-triangles in such a way that the number of triangles incident at any vertex is minimized.
Case 2: The polygon has exactly two internal bounding edges.
In this case, the polygon is an n-sided analogue of a sleeve triangle and ought to constitute the intermediate portion of a limb or a torso. Here we obtain a triangulation of the polygon, consisting of only ***S***-triangles such that the number of triangles incident at a vertex is minimized.
Case 3 : The polygon has m ($2<m\leq n$) internal bounding edges.
In this case, the polygon is an m-sided analogue of a junction triangle. Here, we choose a triangulation that consists of $n-m$ ***S***-triangles and $m-2$ ***J***-triangles. Also, the number of ***J***-triangles adjacent to the bounding edges of the polygon is minimized, and the ***J***-triangles are consolidated to form a convex m-sided subpolygon consisting of only ***J***-triangles. Indeed, such a triangulation of the polygon's interior may be arrived at by starting with any triangulation and changing the morphological status (i.e., ***T***, ***S***, or ***J***) and configurations of pairwise adjacent triangles by flipping their common edge. Thus flipping the common edge between a ***J***-triangle and a ***T***-triangle yields two adjacent ***S***-triangles. Flipping among other possible adjacent pair configurations (eg., ***J-S***, ***S-T***, ***J-J***) does not change the status of the resultant triangle pair but changes their relative configurations. Flipping the common edge of two adjacent co-circular triangles does not violate the definition of the CDT.
The nature of the triangulation within the m-sided subpolygon of ***J***-triangles is irrelevant for the purpose of skeletonization. This is because the part of the skeleton inside this subpolygon will be obtained by joining the midpoints of its bounding edges to its circumcenter (if it lies in the subpolygon) or to the midpoint of the longest bounding edge.
Case 4 : The polygon has no internal bounding edges.
In this case, the polygon is itself the whole polygonized shape, which is a convex polygon with co-circular vertices. If the polygon contains its circumcenter, then it is desirable to minimize the number of ***S***-triangles to 1 (n odd) or 0 (n even), maximize the number of ***T***-triangles to $[n/2]$ and consequently the number of ***J***-triangles to $[n/2]-2$. This interprets the polygonized shape as a discretized circle, with only the circumcenter indicating its skeleton. If the polygon does not contain its circumcenter, then it is desirable to maximize the number of ***S***-triangles and minimize the number of ***J***-triangles to zero (and subsequently the number of ***T***-triangles to two). This is done by selecting the longest edge of the polygon, and for each vertex subtending this edge choosing its nearest neighbor and joining the neighbors of this vertex to form a ***T***-triangle. The remaining $n-4$ triangles are constructed to be ***S***-triangles that minimize the number of triangles incident at any vertex of the polygon.
Such optimizations and constraints as described above may be formally incorporated into the definition of the CDT to disambiguate degenerate cases and adapt its properties to shape analysis.
Using the above-identified structures (i.e., limbs and torsos), we can define the CAT of a polygonized shape.
The internal edges (i.e., the edges connecting nonadjacent boundary points) of the CDT of a polygonized shape are maximal chords (loosely referred to as "chords"). The midpoints of chords along with half their lengths will constitute ordered pairs (p,δ) that constitute the CAT of a polygonized shape.
The vertex opposite the internal edge of a terminal triangle, together with the length of (either) one of the external edges incident upon it will constitute a special ordered pair of the form $(p,-\delta)$-a *terminal ordered*

pair. Here p is the vertex (represented by its pair of coordinates,) and δ is the length of one of the external (boundary) edges incident upon it. The negative sign in front of δ just serves to distinguish this type of ordered pair from the others and has no computational use or significance.
Two ordered pairs corresponding to the adjacent sides of the same triangle are called *adjacent pairs*. Each adjacent pair defines a unique triangle, and a unique adjacent pair defines each ***S***-triangle.
The CAT of a polygonized shape is then a list of sequences of ordered pairs, wherein each sequence characterizes a limb or a torso of the shape in accordance with the following definition.
Definition 5: The chordal axis transform of a polygonized shape is the list of all sequences, of length at least 3, of ordered pairs (p,δ), obtained from the CDT of the shape, of the form
$[(p_1,-\delta_1), (p_2,\delta_2), \ldots, (p_{n-1},\delta_{n-1}), (p_n,\delta_n)\ ]$ (corresponding to a limb)
or of the form
$[(p_1,\delta_1), (p_2,\delta_2), \ldots, (p_{n-1},\delta_{n1}), (p_n,\delta_n)]$ (corresponding to a torso)
such that, in each sequence

1) each ordered pair (p_i,δ_i) ($1 < i < n$) occurs exactly once;
2) (p_i,δ_i) and (p_{i+1},δ_{i+1}) ($1 \le i < n$) are adjacent pairs;
3) no two adjacent pairs define the same triangle, unless that triangle is a ***J***–triangle *and*
$$(p_2,\delta_2) \neq (p_{n-1},\delta_{n-1});$$
4) the adjacent pairs (p_1,δ_1), (p_2,δ_2) and (p_{n-1},δ_{n-1}), (p_n,δ_n) both define ***J***-triangles (which may be the same triangle, as in the case of loops); and
5) $(p_1,-\delta_1)$ is a terminal ordered pair that occurs only at the beginning of a sequence, characterizing a limb.

The above definition of the CAT of a polygonal shape completely characterizes the shape. Indeed, as will be shown later, it is possible to recover the polygonal shape, *exactly*, from its CAT. In fact, not only is it possible to recover the shape completely but also one can reconstruct the CDT of the shape from its CAT. The discrete CAT and its ability to segment a polygonal shape into limbs and torsos is displayed geometrically by joining p_i and p_{i+1} whenever the adjacent pair (p_i,δ_i) and (p_{i+1},δ_{i+1}) define an ***S***-triangle. This results in a disconnected skeletal segmentation of the shape structure by means of polygonal arcs (Figs.6, 12).

3. THE CAT SKELETON

It is possible to construct a one-dimensional retract of the polygonal shape, that has the same connectivity as the shape and serves as the local axis of symmetry of the shape, namely its skeleton (Figs.8, 13). A brief outline of the algorithm for constructing the skeleton of a polygonal shape from its CAT is given in the following steps

1) Join p_i and p_{i+1} by a straight line segment whenever the adjacent pair (p_i,δ_i) and (p_{i+1},δ_{i+1}) define an ***S***-triangle (i.e., whenever $2 \le i \le n-2$)
2) In each of the ***J***-triangles determined by the adjacent pairs (p_1,δ_1), (p_2,δ_2) and (p_{n-1},δ_{n-1}), (p_n,δ_n) of a sequence in the CAT, perform the following:

(a) join the midpoints of the sides of the ***J***-triangle to its circumcenter (the intersection of the perpendicular bisectors of the sides of the triangle) if the triangle is acute (i.e., if the circumcenter lies inside the triangle); and
(b) join the midpoint of the longest side of the triangle to the midpoints of the other two sides if the triangle is not acute (i.e., if the circumcenter lies outside the triangle).

When ***J***-triangles form a convex subpolygon whose vertices lie on the same empty circle, the skeleton inside the subpolygon is obtained by joining the midpoints of its bounding edges to the circumcenter of the empty circle (if it lies inside the subpolygon) or to the midpoint of its longest bounding edge.

To accomplish step 2 above, it is necessary to obtain the vertices of the relevant ***J***-triangles. This can be done using the reconstruction algorithm laid out in the next section.

4. INVERTING THE CAT

The following is a brief outline of the basic steps of the algorithm to recover the boundary of a polygonal shape from its CAT.
For the purposes of inversion, it is only necessary to consider sequences of the CAT that are limbs, or torso sequences that have length greater than three. In other words, only those sequences that contain adjacent

pairs corresponding to sleeve triangles or terminal ordered pairs need to be considered. This is because only terminal and sleeve triangles share edges with the polygonal boundary. In any such sequence, it is sufficient to describe how the *span* (i.e., the subtending boundary points of a maximal chord) of any ordered pair (p_i, δ_i) $(1 < i < n)$ may be recovered. For this purpose, consider the ordered pair (p_i, δ_i) along with its adjacent ordered pairs (p_{i-1}, δ_{i-1}) and (p_{i+1}, δ_{i+1}) (Fig.14). Let Δ_{i-1} be the triangle corresponding to the adjacent pair (p_{i-1}, δ_{i-1}), (p_i, δ_i), and let Δ_i be the triangle corresponding to the adjacent pair (p_i, δ_i), (p_{i+1}, δ_{i+1}). If the distance $d(p_{i-1}, p_i)$ between the points p_{i-1} and p_i is denoted ρ_{i-1}, then half the height of the triangle Δ_{i-1} is given by

$$\theta_{i-1} = \sqrt{(\delta_{i-1} + \delta_i + \rho_{i-1})(\delta_{i-1} + \delta_i - \rho_{i-1})(\delta_{i-1} - \delta_i + \rho_{i-1})(\delta_i - \delta_{i-1} + \rho_{i-1})}\,/2\rho_{i-1}$$

If μ_{i-1} is the slope of the straight line $\mathbf{\Lambda}_{i-1}$ passing through p_{i-1} and p_i, then the two straight lines, $\mathbf{\Lambda}^{+}_{i-1}$ and $\mathbf{\Lambda}^{-}_{i-1}$, parallel to the line $\mathbf{\Lambda}_{i-1}$ and at a distance θ_{i-1} on either side of it, are given by

$$\mathbf{\Lambda}^{\pm}_{i-1} : y = \mu_{i-1} x + C_{i-1} \pm \sigma_y^{i-1} \theta_{i-1} \sqrt{1 + \mu_{i-1}^2}$$

where

$$\mathbf{\Lambda}_{i-1} : y = \mu_{i-1} x + C_{i-1}$$

and, if p_{i-1} and p_i are given by (x_{i-1}, y_{i-1}) and (x_i, y_i) then

$$\mu_{i-1} = \frac{y_i - y_{i-1}}{x_i - x_{i-1}} \quad \text{and} \quad \sigma_y^{i-1} = \text{sgn}\,(y_i - y_{i-1})$$

Similar definitions and equations can be set forth for the adjacent pair (p_i, δ_i), (p_{i+1}, δ_{i+1}).
The intersection of the lines $\mathbf{\Lambda}^{+}_{i-1}$ and $\mathbf{\Lambda}^{+}_{i}$ yields the boundary point b^{+}_{i}, and, similarly, the intersection of the lines $\mathbf{\Lambda}^{-}_{i-1}$ and $\mathbf{\Lambda}^{-}_{i}$ yields the boundary point b^{-}_{i}. The pair b^{+}_{i} , b^{-}_{i} form the span of the ordered pair (p_i, δ_i).
Thus, two sequences of pairwise adjacent boundary vertices are generated:

$$\{ b^{+}_{2}, \ldots, b^{+}_{n-1} \} \text{ and } \{ b^{-}_{2}, \ldots, b^{-}_{n-1} \}$$

At this point, if needed, the CDT can also be recovered by joining the pairs b^{+}_{i} , b^{-}_{i} by a line segment.
Some vertices in each of the above sequences appear a multiplicity of times successively; this is because consecutive intersection points of lines sometimes coincide. This multiplicity is resolved by retaining only one copy of each vertex. In the case wherein the sequence of ordered pairs corresponds to a limb (this is signified by the first ordered pair being a terminal ordered pair represented by $(p_1, -\delta_1)$), the vertex p_1 is the common boundary neighbor of the vertices b^{+}_{2} and b^{-}_{2}. Thus, for a limb (Fig.14) sequence, a single contiguous chain of pairwise adjacent polygonal boundary vertices, given by $\{ b^{-}_{n-1}, \ldots, b^{-}_{2}, p_1, b^{+}_{2}, \ldots, b^{+}_{n-1} \}$, is obtained, whereas for a torso sequence two disjoint contiguous chains of boundary vertices, given by $\{ b^{-}_{n-1}, \ldots, b^{-}_{2} \}$ and $\{ b^{+}_{2}, \ldots, b^{+}_{n-1} \}$, are obtained. When two boundary vertices of two distinct CAT sequences coincide, they are taken to be the same vertex, and the boundary vertex sequences are merged at this vertex. Performing the above constructions for all limb sequences and for all torso sequences of length greater than 3 in the CAT of a polygonal shape, all the boundary vertices of the polygonal shape along with adjacency relations are obtained. When two neighboring adjacent pairs are collinear, then the line pairs $\mathbf{\Lambda}^{+}_{i-1}$, $\mathbf{\Lambda}^{+}_{i}$ and $\mathbf{\Lambda}^{-}_{i-1}$, $\mathbf{\Lambda}^{-}_{i}$ fail to intersect at a point. In this case, a slightly different reconstruction involving intersections of circles with lines is employed for such sections. We will not elaborate on this here, but simply remark that the CAT has all the information required to recover the shape boundary efficiently and exactly.

5. PRUNING OF SHAPE FEATURES AND SKELETONS

When the boundary of a shape is densely sampled to include noise or minor undulations, the skeleton of the resulting approximating polygon contains branches that are not morphologically significant (Fig.15a). The shape and its skeleton can be further pruned to isolate its prominent morphological features. This technique removes branches by eliminating minor local boundary features in the CAT where the features do not contribute significantly to the characterization of the overall shape. In the CDT of an approximating polygon, each side of a ***J***-triangle that connects boundary points of the same boundary component subtends

a chain of polygonal vertices that does not include the vertex of the ***J***-triangle opposite to this side (Fig.16). The ratio of morphological significance $\rho = d / |AB|$, of the distance d between the farthest point p of the chain from the side AB, of the junction triangle ABC, is a quantitative indication of the importance of the portion $AopqrsBA$ in describing the overall shape (Fig.16). Whenever a part of a shape (subtended by an edge of a ***J***-triangle) is morphologically insignificant, i.e., whenever ρ is less than some threshold, the part is excised from the shape. The edge subtending the excised part becomes part of the new polygonal boundary (i.e., A and B become neighboring boundary vertices of the modified polygon), while the ***J***-triangle to which the edge belongs becomes an ***S***-triangle. This results in a simplified shape that still represents the salient features of the original shape. Accordingly, the new shape's skeleton does not reflect the morphologically insignificant branches associated with the excised part of the shape (Figs. 15b, 15c, 15d). This process of pruning a shape can also be incorporated directly into the CAT of the shape, since the CAT contains all the information about the shape. The act of pruning does not affect the exact reconstructability of the rest of the shape from its modified CAT.

6. FEATURE-BASED SHAPE REPRESENTATION VIA ATTRIBUTED GRAPHS

The essential structural features and characteristics of a polygonized shape may be abstracted and represented in a coordinate independent manner by means of edge-attributed planar graphs. For a polygonized shape **P**, let CDT(**P**) denote the set of all triangles in its CDT. We define the edge-attributed planar graph **G***w*(**P**) (Fig.17) associated with **P**, with node set ***V*** and edge set ***E***, as follows:

Let each node in ***V*** correspond uniquely to a ***T***-triangle or a ***J***-triangle in CDT(**P**). Conversely, let each ***J***-triangle and each ***T***-triangle in CDT(**P**) correspond uniquely to a node in ***V***. Thus, there is a one-to-one correspondence between the nodes of ***V*** and the set of junction and terminal triangles of CDT(**P**); the size, $|V|$, of ***V*** is $(\Delta_J + \Delta_T)$.

If ***TS…SJ*** is a limb in the CAT of **P**, then an edge in ***E*** joining the two nodes in ***V*** corresponding to the terminal triangle ***T*** and the junction triangle ***J*** of the limb is introduced. Similarly, if ***JS…SJ*** is a torso in the CAT of **P**, then an edge in ***E*** joining the two nodes in ***V*** corresponding to the two junction triangles ***J*** of the torso is introduced. The size, $|E|$, of ***E*** is $(N(\lambda)+N(\tau))$. The resulting graph is planar since it has the same topology as the skeleton of **P**.

Each edge of this graph is attributed with a weight vector, signifying metrical attributes such as length, width, area, etc., of the corresponding limb or torso. The following is an algorithm to obtain the edge-attributed planar graph representation of a polygonized shape from its CDT:

If $[(p_1,\pm\delta_1), (p_2,\delta_2),\ldots, (p_{n-1},\delta_{n-1}), (p_n,\delta_n)\ ]$ is a torso or limb sequence in the CAT of P, the length l of this limb is taken to be

$$l = \sum_{i=2}^{n-2} \rho_i$$

where $\rho_i = d(p_{i-1},p_i)$, as before.

The width w of the limb is taken to be

$$w = 2\left(\frac{\sum_{i=2}^{n-2} \theta_i \rho_i}{\sum_{i=2}^{n-2} \rho_i}\right)$$

where θ_i is, as before, half the height of the triangle defined by the adjacent pair (p_i,δ_i), (p_{i+1},δ_{i+1}).

The variance of the heights is taken as another attribute of the limb to indicate the variation in the width of the limb along its length:

$$v = 4\left(\frac{\sum_{i=2}^{n-2} (\theta_i - w)^2 \rho_i}{\sum_{i=2}^{n-2} \rho_i}\right)$$

The weight vector ϖ for the edge joining the nodes corresponding to the ***T*** and ***J*** triangles of the above limb ***TS…SJ*** is then taken to be $\varpi = (l, w, v)$. A similar computation of the weight vector is performed for

graph edges representing torsos. For limbs and torsos that do not have sleeve triangles (i.e., the limbs and torsos have only two triangles in them) the width of the common edge between the triangles is taken to be the width $\boldsymbol{w}$ of the member. The variance $\boldsymbol{v}$ is set to zero. The length $\boldsymbol{l}$ is taken to be the distance between the midpoint of the edge of the $\boldsymbol{J}$-triangle to the vertex of the $\boldsymbol{T}$-triangle for a limb, and zero for a torso. The individual components of all the weights of the graph $\mathbf{G}_{\boldsymbol{w}}(\mathbf{P})$ may then be separately or jointly normalized with respect to the largest weight in the graph for a scale-invariant representation. The number of components of the weight vector may be increased to include other higher-order statistics or other properties of limbs and torsos, as the application at hand might dictate.

This representation of shapes is invariant under a wide range of transformations on the shape itself; in particular the representation is invariant under affine transformations such as translation, rotation, scaling, and skewing. Also, the size of the graph representation depends only on the structural complexity (i.e., the number of limbs, torsos, junctions, and terminations) of the shape, and not on the particular mode of representation of the shape itself. This is a very important feature in terms of achieving highly economical structural representation and encoding of shape information.

This representation, along with the enumeration of graph components, may then be used for matching and comparing shapes at a high-level using attributed graph matching algorithms.[8] This characterization has several important applications in computer vision and image understanding. The problem of planar graph isomorphism has been solved efficiently. For shape recognition purposes, it is necessary, though, to be able to compare graphs that have significant similarities without being isomorphic. However, this more difficult problem of sub-graph isomorphism is still very much open. In a subsequent paper,[4] presented at this conference, we present a feature-based linguistic encoding of shapes via context-free grammars, that enables hierarchical matching of shapes for the purposes of object recognition.

7. CONCLUSIONS

We have presented a novel and general geometric filtering scheme for efficiently extracting salient morphological features from shapes and representing shapes at a high level in terms of their features. The flexibility of this scheme allows for many variations and adaptations to suit specific applications. The algorithms developed and used in this scheme are simple and highly efficient in their time and space complexities. They have been implemented to perform in real time with video imagery.[6] The methods described here have important applications in computer vision, image understanding, pattern analysis, and artificial intelligence. This work is part of an ongoing effort to automate the intelligent assessment of image data for surveillance and remote sensing applications.

8. ACKNOWLEDGMENTS

This work is supported by DOE LDRD-ER (2000021). Related applied research and implementation in the area of video surveillance, is supported by DOE NN-20.

REFERENCES

1. H. Blum, "A Transformation for Extracting New Descriptors of Shape," *Symposium on Models for Speech and Visual Form*, Weiant Whaten-Dunn (Ed.) MIT Press, Cambridge, MA, 1967.
2. L. Prasad, "Morphological Analysis of Shapes," CNLS Newsletter, No. 139, July 1997, LALP-97-010-139, Center for Nonlinear Studies, Los Alamos National Laboratory.
3. L. Prasad, R. L. Rao "Multi-scale Discretization of Shape Contours," in Mathematical Imaging, **4117** Vision Geometry IX, *Proc. of the 45th SPIE Annual Meeting*, San Diego, USA, 2000.
4. L. Prasad, A. N. Skourikhine, B. R. Schlei "Feature-based Syntactic and Metric Shape Recognition," in Mathematical Imaging, **4117** Vision Geometry IX, *Proc. of the 45th SPIE Annual Meeting*, San Diego, USA, 2000.
5. B. R. Schlei, L. Prasad, A. N. Skourikhine, "Geometric Morphology of Granular Materials," in Mathematical Imaging, **4117** Vision Geometry IX, *Proc. of the 45th SPIE Annual Meeting*, San Diego, USA, 2000.
6. B. R. Schlei, "GEOFILT-Geometric Filtering Code, Version 1.0," *Los Alamos Computer Code* LA-CC-00-31, Los Alamos National Laboratory, 2000; (www.nis.lanl.gov/~bschlei/eprint.html.)
7. R.L. Ogniewicz, "Skeleton-space: A multiscale shape description combining region and boundary information," in Proc. IEEE Conf. on Computer Vision and Pattern Recognition, Seattle, WA, pp. 746-751, June 1994.

8. S. Gold, A. Rangarajan, "A graduated assignment algorithm for graph matching," IEEE Trans. Pattern Anal. Mach. Intell. 18(4), 377-388 (1996).

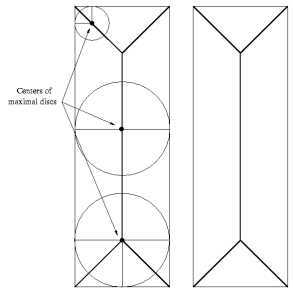

Fig.1. A rectangle's MAT.

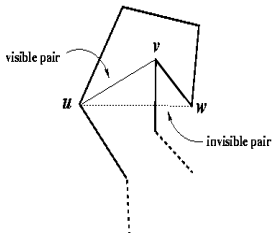

Fig.2. A visible vertex pair.

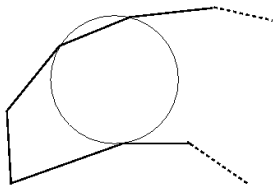

Fig.3. An empty circle.

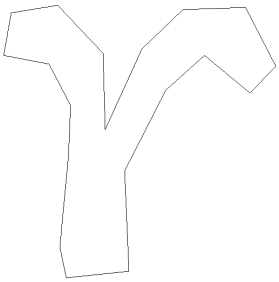

Fig.4. Polygonized shape of handwritten letter "***r***".

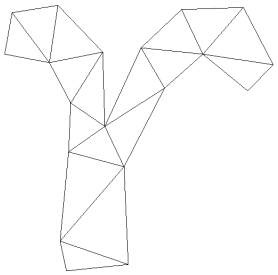

Fig.5. CDT.

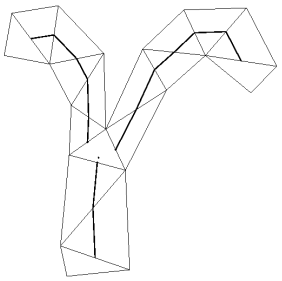

Fig.6. Segmentation induced by CAT.

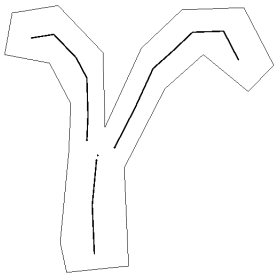

Fig.7. CAT protoskeleton.

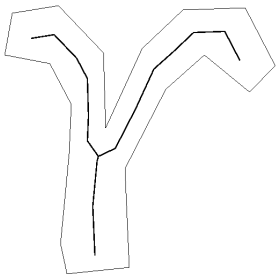

Fig.8. CAT skeleton.

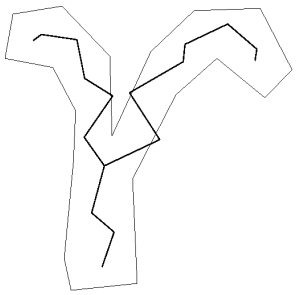

Fig.9. Discrete MAT skeleton.

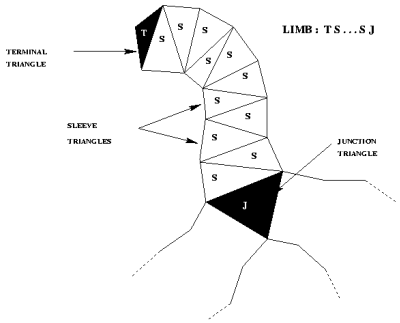

Fig.10. A limb chain complex.

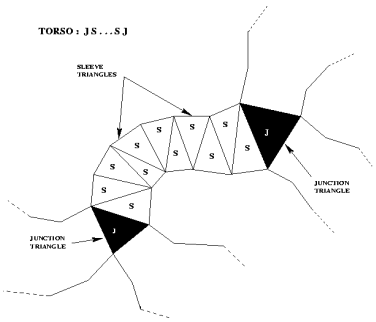

Fig.11. A torso chain complex.

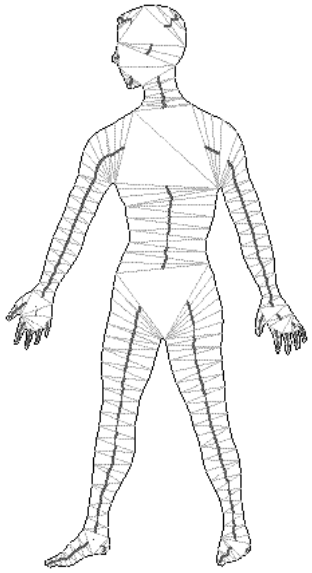

Fig.12. Segmentation induced by the CAT.

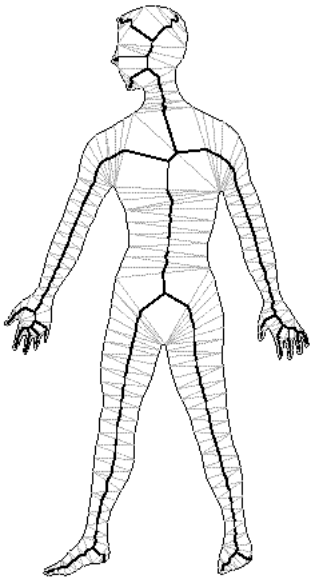

Fig.13. Skeleton induced by the CAT.

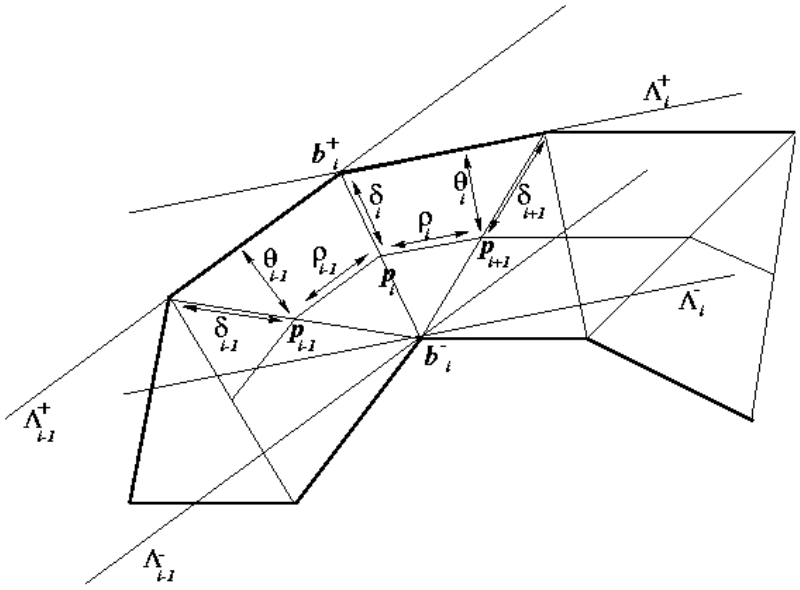

Fig.14. Reconstruction of the span (b^+_i, b^-_i) of an ordered pair (p_i, δ_i) from the CAT.

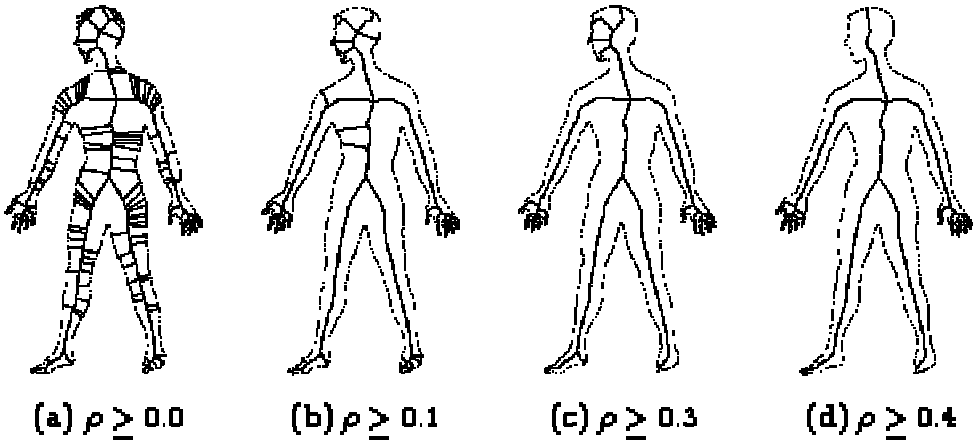

Fig.15. Pruning of a densely sampled human shape with noisy boundary, for various values of the threshold ρ, with resulting simplified skeletons. In each case the shape components whose ratio of morphological significance exceed ρ are retained.

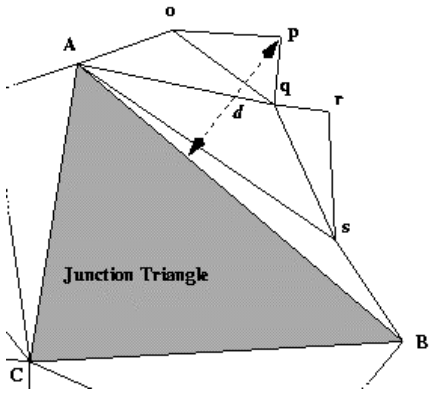

Fig.16. Pruning a morphologically insignificant feature; A&B must belong to the same contour component.

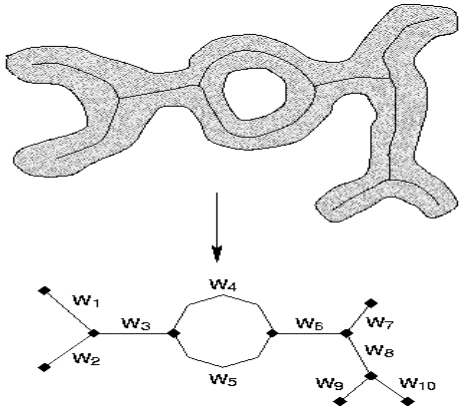

Fig.17. A shape with its skeleton shown, and its edge-attributed planar graph with weights $\mathbf{W_i}$.